arXiv:1404.3248v2 [cs.DS] 21 Jan 2015

Solving Optimization Problems with Diseconomies of Scale via Decoupling*

Konstantin Makarychev
Microsoft Research

Maxim Sviridenko
Yahoo! Labs

Abstract

We present a new framework for solving optimization problems with a diseconomy of scale. In such problems, our goal is to minimize the cost of resources used to perform a certain task. The cost of resources grows superlinearly, as x^q, $q \geq 1$, with the amount x of resources used. We define a novel linear programming relaxation for such problems, and then show that the integrality gap of the relaxation is A_q, where A_q is the q-th moment of the Poisson random variable with parameter 1. Using our framework, we obtain approximation algorithms for the Minimum Energy Efficient Routing, Minimum Degree Balanced Spanning Tree, Load Balancing on Unrelated Parallel Machines, and Unrelated Parallel Machine Scheduling with Nonlinear Functions of Completion Times problems.

Our analysis relies on the decoupling inequality for nonnegative random variables. The inequality states that

$$\|\sum_{i=1}^{n} X_i\|_q \leq C_q \|\sum_{i=1}^{n} Y_i\|_q,$$

where X_i are independent nonnegative random variables, Y_i are possibly dependent nonnegative random variable, and each Y_i has the same distribution as X_i. The inequality was proved by de la Peña in 1990. De la Peña, Ibragimov, and Sharakhmetov showed that $C_q \leq 2$ for $q \in (1,2)$ and $C_q \leq A_q^{1/q}$ for $q \geq 2$. We show that the optimal constant is $C_q = A_q^{1/q}$ for any $q \geq 1$. We then prove a more general inequality: For every convex function φ,

$$\mathbb{E}\left[\varphi\Big(\sum_{i=1}^{n} X_i\Big)\right] \leq \mathbb{E}\left[\varphi\Big(P\sum_{i=1}^{n} Y_i\Big)\right],$$

and, for every *concave* function ψ,

$$\mathbb{E}\left[\psi\Big(\sum_{i=1}^{n} X_i\Big)\right] \geq \mathbb{E}\left[\psi\Big(P\sum_{i=1}^{n} Y_i\Big)\right],$$

where P is a Poisson random variable with parameter 1 independent of the random variables Y_i.

*The conference version of this paper appeared at FOCS 2014.

1 Introduction

In this paper, we study combinatorial optimization problems with a diseconomy of scale. We consider problems in which we need to minimize the cost of resources used to accomplish a certain task. Often, the cost grows linearly with the amount of resources used. In some applications, the cost is sublinear e.g., if we can get a discount when we buy resources in bulk. Such phenomenon is known as "economy of scale". However, in many applications the cost is superlinear. In such cases, we say that the cost function exhibits a "diseconomy of scale". A good example of a diseconomy of scale is the cost of energy used for computing. Modern hardware can run at different processing speeds. As we increase the speed, the energy consumption grows superlinearly. It can be modeled as a function $P(s) = cs^q$ of the processing speed s, where c and q are parameters that depend on the specific hardware. Typically, $q \in (1, 3]$ (see e.g., [2, 20, 39]).

As a running example, consider the Minimum Power Routing problem studied by Andrews, Fernández Anta, Zhang, and Zhao [3]. We are given a graph $G = (V, E)$ and a set of demands $\mathcal{D} = \{(d_i, s_i, t_i)\}$. Our goal is to route d_i ($d_i \in \mathbb{N}$) units of demand i from the source $s_i \in V$ to the destination $t_i \in V$ such that every demand i is routed along a single path p_i (i.e. we need to find an unsplittable multi-commodity flow). We want to minimize the energy cost. Every link (edge) $e \in E$ uses $f_e(x_e) = c_e x_e^q$ units of power, where c_e is a scaling parameter depending on the link e, and x_e is the load on e.

The straightforward approach to solving this problem is as follows. We define a mathematical programming relaxation that routes demands fractionally. It sends $y_{i,p}d_i$ units of demand via the path p connecting s_i to t_i. We require that $\sum_p y_{i,p} = 1$ for every demand i. The objective function is to minimize

$$\min \sum_{e \in E} c_e x_e^q = \min \sum_{e \in E} c_e \Big(\sum_{p: e \in p} y_{i,p} d_i \Big)^q,$$

where $x_e = \sum_{p: e \in p} y_{i,p} d_i$ is the load on the link e. This relaxation can be solved in polynomial time, since the objective function is convex (for $q \geq 1$). But, unfortunately, the integrality gap of this relaxation is $\Omega(n^{q-1})$ [3]. Andrews et al. [3] gave the following integrality gap example. Consider two vertices s and t connected via n disjoint paths. Our goal is to route 1 unit of flow integrally from s to t. The optimal solution pays 1. The LP may cheat by routing $1/n$ units of flow via n disjoint paths. Then, it pays only $n \times (1/n)^q = n^{1-q}$.

For the case of uniform demands, i.e., for the case when all $d_i = d$, Andrews et al. [3] suggested a different objective function:

$$\min \sum_{e \in E} c_e \max\{x_e^q, d^{q-1} x_e\}.$$

The objective function is valid, because in the integral case, x_e must be a multiple of d, and thus $x_e^q \geq d^{q-1} x_e$. Andrews et al. [3] proved that the integrality gap of this relaxation is a constant. Bampis et al. [9] improved the bound to the *fractional Bell number* A_q that is defined as follows: A_q is the q-th moment of the Poisson random variable P_1 with parameter 1 (see Figure 2 in Appendix H). I.e.,

$$A_q = \mathbb{E}\left[P_1^q\right] = \sum_{t=1}^{+\infty} t^q \frac{e^{-1}}{t!}. \tag{1}$$

For the case of general demands no constant approximation was known. The best known approximation due to Andrews et al. [3] was $O(k + \log^{q-1} \Delta)$ where $k = |\mathcal{D}|$ is the number of demands and $\Delta = \max_i d_i$ is the size of the largest demand (Theorem 8 in [3]).

In this work, we give an A_q-approximation algorithm for the general case and thus close the gap between the case of uniform and non-uniform demands. Our approximation algorithm uses

a general framework for solving problems with a diseconomy of scale which we present in this paper. We use this framework to obtain approximation algorithms for several other combinatorial optimization problems. We give $A_q^{1/q}$-approximation algorithm for Load Balancing on Unrelated Parallel Machines (see Section 2.2), $2^q A_q$-approximation algorithm for Unrelated Parallel Machine Scheduling with Non-linear Functions of Completion Times (see Section 2.3) and A_q-approximation algorithm for the Minimum Degree Balanced Spanning Tree problem (see Section 2.4). The best previously known bound for the first problem with $q \in [1, 2]$ was $2^{1/q}$ (see Figure 3 for comparison). The bound is due to Kumar, Marathe, Parthasarathy and Srinivasan [22]. There were no known approximation guarantees for the latter problems.

In the analysis, we use the de la Peña decoupling inequality [26, 27].

Theorem 1.1 (de la Peña [26, 27])**.** *Let $Y_1, \ldots, Y_n$ be jointly distributed nonnegative (non-independent) random variables, and let $X_1, \ldots, X_n$ be independent random variables such that each X_i has the same distribution as Y_i. Then, for every $q \geq 1$,*

$$\big\| \sum_{i=1}^{n} X_i \big\|_q \leq C_q \, \big\| \sum_{i=1}^{n} Y_i \big\|_q, \tag{2}$$

for some universal constant C_q.

De la Peña, Ibragimov, and Sharakhmetov ([28], Corollary 3.4) showed that $C_q \leq 2$ for $q \in [1, 2]$, and $C_q \leq A_q^{1/q}$ for $q \geq 2$. We give an alternative proof of this inequality, and show that the inequality holds for $C_q = A_q^{1/q}$ for any $q \geq 1$, and moreover this bound is tight (for any $q \geq 1$). Thus, we improve the known upper bound for C_q for $q \in (1, 2)$.

Theorem 1.2. *Inequality (2) holds for $C_q = A_q^{1/q}$, where A_q is the fractional Bell number (see Equation (1) and Figure 2). Moreover, $A_q^{1/q}$ is the optimal upper bound on C_q.*

In fact we prove a more general inequality for arbitrary convex functions and an analogous inequality for concave functions. In Section 6 (see Corollary 6.2), we extend this theorem to negatively associated random variables X_i.

1.1 General Framework

We now describe the general framework for solving problems with a diseconomy of scale. We consider optimization problems with n decision variables $y_1, \ldots, y_n \in \{0, 1\}$. We assume that the objective function equals the sum of k terms, where the j-th term is of the form

$$f_j\Big(\sum_{i=1}^{n} d_{ij} y_i \Big),$$

here f_j's are nonnegative monotonically nondecreasing convex functions, $d_{ij} \geq 0$ are parameters. The vector $y = (y_1, \ldots, y_n)$ must satisfy the constraint $y \in \mathcal{P}$ for some polytope $\mathcal{P} \subset [0, 1]^n$. Therefore, the optimization problem can be written as the following boolean convex program (IP):

$$\min \quad \sum_{j \in [k]} f_j\Big(\sum_{i \in [n]} d_{ij} y_i \Big) \tag{3}$$

$$y \in \mathcal{P} \tag{4}$$

$$y \in \{0, 1\}^n \tag{5}$$

We assume that we can optimize any linear function over the polytope $\mathcal{P}$ in polynomial time (e.g., $\mathcal{P}$ is defined by polynomially many linear inequalities, or there exists a separation oracle for $\mathcal{P}$). Thus, if we replace the integrality constraint (5) with the relaxed constraint $y \in [0,1]^n$ (which is redundant, since $\mathcal{P} \subset [0,1]^n$), we will get a convex programming problem that can be solved in polynomial time (see [11]). However, as we have seen in the example of Minimum Power Routing, the integrality gap of the relaxation can be as large as $\Omega(n^{q-1})$ for $f_j(t) = t^q$.

In this work, we introduce a linear programming relaxation of (3)-(5) that has an integrality gap of A_q for $f_j(t) = c_j t^q$ under certain assumptions on the polytope $\mathcal{P}$. We define auxiliary variables z_{jS} for all $S \subset [n]$ and $j \in [k]$. In the integral solution, $z_{jS} = 1$ if and only if $y_i = 1$ for $i \in S$ and $y_i = 0$ for $i \notin S$.

$$\min \sum_{j \in [k]} \sum_{S \subseteq [n]} f_j \left(\sum_{i \in S} d_{ij} \right) z_{jS} \tag{6}$$

$$y \in \mathcal{P}, \tag{7}$$

$$\sum_{S \subseteq [n]} z_{jS} = 1, \qquad \forall j \in [k] \tag{8}$$

$$\sum_{S : i \in S} z_{jS} = y_i, \qquad \forall i \in [n], j \in [k] \tag{9}$$

$$z_{jS} \geq 0, \qquad \forall S \subseteq [n], j \in [k] \tag{10}$$

Remark 1.3. *In the integral solution, $z_{j'S} = z_{j''S}$ for all j' and j''. The reason why we introduced many copies of the same integral variable z_S to the LP is that the LP above is easier to solve than the LP with an extra constraint $z_{j'S} = z_{j''S}$.*

Optimization problem (6)-(10) is a relaxation of the original problem (3)-(5). The LP has exponentially many variables. We show, however, that the optimal solution to this LP can be found in polynomial time up to an arbitrary accuracy $(1+\varepsilon)$.

We shall assume that all d_{ij} are integral and polynomially bounded, or, more generally, that all d_{ij} are multiples of some δ_j and that d_{ij}/δ_j are polynomially bounded in n and $1/\varepsilon$. Given an arbitrary instance of the problem, it is easy to round all d_{ij}'s to multiples of a sufficiently small δ_j, so that the cost of any solution changes by at most $(1+\varepsilon)$ assuming that functions f_j satisfy some mild conditions. In Section D (Theorem D.1), we show how to pick δ_j for functions f_j satisfying the following conditions:

1. Each f_j is a convex increasing function;
2. For each j, $f_j(0) = 0$;
3. For each j and $t \in (0, \sum_i d_{ij}]$, $t(\log f_j(t))' \leq p(n)$, where p is some polynomial.

Note, that for $f(t) = ct^q$, we have $t(\log f(t))' = t(\log ct^q)' = t \cdot (q/t) = q$. So functions $f_j = c_j t^{q_j}$ satisfy the conditions of Theorem D.1 if q_j are polynomially bounded.

Theorem 1.4. *Suppose that there exists a polynomial time separation oracle for the polytope $\mathcal{P}$, $f_j(t)$ is computable in polynomial time as a function of j and t, and all d_{ij} are multiples of δ_j such that d_{ij}/δ_j are polynomially bounded in n and $1/\varepsilon$. Then, for every $\varepsilon > 0$, there exists a polynomial time algorithm that finds a $(1+\varepsilon)$-approximately optimal solution to LP (6)-(10).*

For a convex non-decreasing function f define $A(f)$ as follows:

$$A(f) = \sup_{t>0} \mathbb{E}\Big[\frac{f(tP)}{f(t)}\Big], \tag{11}$$

where P is a Poisson random variable with parameter 1. Note that

$$A(ct^q) = \sup_{t>0} \mathbb{E}\Big[\frac{c(tP)^q}{ct^q}\Big] = \mathbb{E}[P^q] = A_q.$$

We prove the following theorem.

Theorem 1.5. *Let $D_j = \{i : d_{ij} \neq 0\}$. Assume that there exists a randomized algorithm R that given a $y \in \mathcal{P}$, returns a random integral point $R(y)$ in $\mathcal{P} \cap \{0,1\}^n$ such that*

1. *$\Pr(R_i(y) = 1) = y_i$ for all i (where $R_i(y)$ is the i-th coordinate of $R(y)$);*
2. *Random variables $\{R_i(y)\}_{i \in D_j}$ are independent or negatively associated (see Section 6 for the definition) for every j.*

Then, for every feasible solution (y^*, z^*) to LP (6)-(10), we have

$$\mathbb{E}\Big[\sum_{j\in[k]} f_j\Big(\sum_{i\in[n]} d_{ij} R_i(y^*)\Big)\Big] \leq \sum_{j\in[k]} A(f_j) \sum_{S\subseteq[n]} f_j\Big(\sum_{i\in S} d_{ij}\Big) z^*_{jS}, \tag{12}$$

where $A(f_j)$ is defined as in (11). Particularly, since LP (6)-(10) is a relaxation for IP (3)-(5), if (y^*, z^*) is a $(1+\varepsilon)$-approximately optimal solution to LP (6)-(10), then

$$\mathbb{E}\Big[\sum_{j\in[k]} f_j\Big(\sum_{i\in[n]} d_{ij} R_i(y^*)\Big)\Big] \leq (1+\varepsilon) \max_j(A(f_j))\, IP,$$

where IP is the optimal cost of the boolean convex program (3)-(5).

This theorem guarantees that an algorithm R satisfying conditions (1) and (2) has an approximation ratio of $(1+\varepsilon)A_q$ for $f_j(t) = c_j t^q$.

In the next section, Section 2, we show how to use the framework to obtain approximation algorithms for four different combinatorial optimization problems. Then, in Section 3, we give an efficient algorithm for solving LP (6)-(10). In Section 4, we prove the main theorem – Theorem 1.5. The proof easily follows from the decoupling inequality, which we prove in Section 5. Finally, in Section 7, we describe some generalizations of our framework.

2 Applications

In this section, we show applications of our general technique. We start with the problem discussed in the introduction – Energy Efficient Routing. Recall, that Andrews et al. [3] gave an $O(k+\log^{q-1}\Delta)$-approximation algorithm for this problem where $k = |\mathcal{D}|$ and $\Delta = \max_{i\in\mathcal{D}} d_i$ (Theorem 8 in [3]). We give an $(1+\varepsilon)A_q$-approximation algorithm for any fixed $\varepsilon > 0$.

2.1 Energy Efficient Routing

We write a standard integer program. Each variable $y_{i,e} \in \{0,1\}$ indicates whether the edge e is used to route the flow from s_i to t_i. Below, $\Gamma^+(u)$ denotes the set of edges outgoing from u; $\Gamma^-(u)$ denotes the set of edges incoming to u.

$$\min \sum_{e \in E} c_e \left(\sum_{i \in \mathcal{D}} d_i y_{i,e} \right)^{q_e} \tag{13}$$

$$\sum_{e \in \Gamma^+(u)} y_{i,e} = \sum_{e \in \Gamma^-(u)} y_{i,e}, \qquad \forall i,\ u \in V \setminus \{s_i, t_i\} \tag{14}$$

$$\sum_{e \in \Gamma^+(s_i)} y_{i,e} = 1, \qquad \forall i \tag{15}$$

$$\sum_{e \in \Gamma^-(t_i)} y_{i,e} = 1, \qquad \forall i \tag{16}$$

$$y_{i,e} \in \{0,1\}, \qquad \forall i,\ e \in E \tag{17}$$

Using Theorem 1.4, we obtain an almost optimal fractional solution (y, z) of LP relaxation (6)-(10) of IP (13)-(17). We apply randomized rounding in order to select a path for each demand. Specifically, for each demand $i \in \mathcal{D}$, we consider the standard flow decomposition into paths: In the decomposition, each path p connecting s_i to t_i has a weight $\lambda_{i,p} \in \mathbb{R}^+$. For every edge e, $\sum_{p: e \in p} \lambda_{i,p} = d_i y_{i,e}$; and $\sum_p \lambda_{i,p} = d_i$. For each i, the approximation algorithm picks one path p connecting s_i to t_i at random with probability $\lambda_{i,p}/d_i$, and routes all demands from s_i to p_i via p. Thus, the algorithm always obtains a feasible solution.

We verify that the integral solution corresponding to this combinatorial solution satisfies the conditions of Theorem 1.5. Let $R_{i,e}(y)$ be the integral solution, i.e., let $R_{i,e}(y) = 1$ if the edge e is chosen in the path connecting s_i and t_i. First, $R_{i,e}(y) = 1$ if the path connecting s_i and t_i contains e, thus

$$\Pr(R_{i,e}(y) = 1) = \sum_{p: e \in p} \lambda_{i,p}/d_i = y_{i,e}.$$

Second, the paths for all demands are chosen independently. Each $R_{i,e}(y)$ depends only on paths that connect s_i to t_i. Thus all random variables $R_{i,e}(y)$ (for a fixed e) are independent. Therefore, by Theorem 1.5, the cost of the solution obtained by the algorithm is bounded by $(1+\varepsilon)A_q\, OPT$, where OPT is the cost of the optimal solution to the integer program which is exactly equivalent to the Minimum Energy Efficient Routing problem.

2.2 Load Balancing on Unrelated Parallel Machines

We are given n jobs and m machines. The processing time of the job $j \in [n]$ assigned to the machine $i \in [m]$ is $p_{ij} \geq 0$. The goal is to assign jobs to machines to minimize the ℓ_q-norm of machines loads. Formally, we partition the set of jobs into m sets $S_1, \ldots, S_m$ to minimize $\left(\sum_{i \in [m]} (\sum_{j \in S_i} p_{ij})^q \right)^{1/q}$. This is a classical scheduling problem which is used to model load balancing in practice[1]. It was previously studied by Azar and Epstein [6] and by Kumar, Marathe, Parthasarathy and Srinivasan [22].

[1] A slight modification of the problem, where the objective is $\min \sum_{i \in [m]} (\sum_{j \in S_i} p_{ij})^q$, can be used for energy efficient scheduling. Imaging that we need to assign n jobs to m processors/cores so that all jobs are completed by a certain deadline D. We can run processors at different speeds s_i. To meet the deadlines we must set $s_i = D^{-1} \sum_{j \in S_i} p_{ij}$. The total power consumption is proportional to $D \times \sum_{i=1}^m s_i^q = D^{1-q} \times \sum_{i=1}^m \left(\sum_{j \in S_i} p_{ij} \right)^q$. For this problem, our algorithm gives $(1+\varepsilon)A_q$ approximation.

Particular, for $q \in (1, 2]$ the best known approximation algorithm has performance guarantee $2^{1/q}$ [22] (Theorem 4.4). We give $\sqrt[q]{(1+\varepsilon)A_q}$-approximation algorithm for any $\varepsilon > 0$ substantially improving upon previous results (see Figure 3).

We formulate the unrelated parallel machine scheduling problem as a boolean nonlinear program:

$$\min \sum_{i\in[m]} \Big(\sum_{j\in[n]} p_{ij}x_{ij}\Big)^q \tag{18}$$

$$\sum_{i\in[m]} x_{ij} = 1, \qquad \forall j \in [n] \tag{19}$$

$$x_{ij} \in \{0, 1\}, \qquad \forall i \in [m],\, j \in [n] \tag{20}$$

Using Theorem 1.4, we obtain an almost optimal fractional solution (x, z) of the LP relaxation (6)–(10) corresponding to the IP (18)–(20). We use the straightforward randomized rounding: we assign each job j to machine i with probability x_{ij}. We claim that, by Theorem 1.5, the expected cost of our integral solution is upper bounded by A_q times the value of the fractional solution (x, z). Indeed, the probability that we assign a job j to machine i is exactly equal to x_{ij}; and we assign job j to machine i independently of other jobs. That implies that our approximation algorithm has a performance guarantee of $\sqrt[q]{(1+\varepsilon)A_q}$ for the ℓ_q-norm objective.

2.3 Unrelated Parallel Machine Scheduling with Nonlinear Functions of Completion Times

As in the previous problem, in Unrelated Parallel Machine Scheduling with Nonlinear Functions of Completion Times, we are given n jobs and m machines. The processing time of the job $j \in [n]$ assigned to the machine $i \in [m]$ is $p_{ij} \geq 0$. We need to assign jobs to machines and set their start times such that job processing intervals do not overlap. The goal is to minimize $\sum_{s=1}^{n} w_j C_j^p$ where C_j is the completion time of job j in the schedule and $p \geq 1$. Using classical scheduling notation this problem can be denoted as $R||\sum_j w_j C_j^p$.

The problem $R||\sum_j w_j C_j^p$ is well studied for $p = 1$. It is known to be APX-hard [19] while the best known approximation algorithm has a performance guarantee of $3/2$ [34, 36]. For $p > 1$ even the single machine scheduling problem is not understood: It is an open problem whether $1||\sum_j w_j C_j^p$ is $\mathcal{NP}$-hard for $p > 0$, $p \neq 1$. Bansal and Pruhs [10] and Stiller and Wiese [38] gave constant factor approximation algorithms for more general functions of completion times for a single machine. However, there were no known approximation algorithms for multiple machines. We show how to use our framework for this problem in Appendix B. Our algorithm gives $2^p A_p$ approximation.

2.4 Degree Balanced Spanning Tree Problem

We are given an undirected graph $G = (V, E)$ with edge weights $w_e \geq 0$. The goal is to find a spanning tree T minimizing the objective function

$$f(T) = \sum_{v\in V} \left(\sum_{e\in\delta(v)\cap T} w_e \right)^q, \tag{21}$$

where $\delta(v)$ is the set of edges in E incident to the vertex v. For $q = 2$, a more general problem was considered before in the Operations Research literature [5, 23, 25, 29] under the name of Adjacent Only Quadratic Spanning Tree Problem. A related problem, known as Degree Bounded Spanning

Tree, received a lot of attention in Theoretical Computer Science [35, 16]. We are not aware of any previous work on Degree Balanced Spanning Tree Problem.

Let x_e be a boolean decision variable such that $x_e = 1$ if we choose edge $e \in E$ to be in our solution (tree) T. We formulate our problem as the following convex boolean optimization problem

$$\begin{aligned} \min \sum_{v \in V} \left(\sum_{e \in \delta(v)} w_e x_e \right)^q & \\ x & \in \mathcal{B}(\mathcal{M}) \\ x_e & \in \{0,1\}, \qquad \forall e \in E, \end{aligned}$$

where $\mathcal{B}(\mathcal{M})$ is the base polymatroid polytope of the graphic matroid in graph G. We refer the reader to Schrijver's book [31] for the definition of the matroid. Using Theorem 1.4, we obtain an almost optimal fractional solution x^* of LP relaxation (6)-(10) corresponding to the above integer problem.

Following Calinescu et al. [7], we define the continuous extension of the objective function (21) for any fractional solution x'

$$F(x') = \sum_{S \subseteq [n]} f(S) \prod_{e \in S} x'_e \prod_{e \notin S} (1 - x'_e),$$

i.e. $F(x')$ is equal to the expected value of the objective function (21) for the set of edges sampled independently at random with probabilities $x'_e, e \in E$. The function F can be approximated with arbitrary polynomially small precision efficiently via sampling (see [7]). By Theorem 1.5, we get the bound $F(x^*) \leq A_q \cdot LP^*$, where LP^* is the value of the LP relaxation (6)–(10) on the fractional solution x^*.

The rounding phase of the algorithm implements the pipage rounding technique [1] adopted to polymatroid polytopes by Calinescu et al. [7]. Calinescu et al. [7] showed that given a matroid $\mathcal{M}$ and a fractional solution $x \in \mathcal{B}(\mathcal{M})$, one can efficiently find two elements, or two edges in our case, e' and e'' such that the new fractional solution $\tilde{x}(\varepsilon)$ defined as $\tilde{x}_{e'}(\varepsilon) = x_{e'} + \varepsilon$, $\tilde{x}_{e''}(\varepsilon) = x_{e''} - \varepsilon$ and $\tilde{x}_e(\varepsilon) = x_e$ for $e \notin \{e', e''\}$ is feasible in the base polymatroid polytope for small positive and for small negative values of ε.

They also showed that if the objective function $f(S)$ is submodular then the function of one variable $F(\tilde{x}(\varepsilon))$ is convex. In our case, the objective function $f(S)$ is supermodular which follows from a more general folklore statement.

Fact 2.1. *The function $f(S) = g(\sum_{i \in S} w_i)$ is supermodular if $w_i \geq 0$ for $i \in [n]$ and $g(x)$ is a convex function of one variable.*

Therefore, the function $F(\tilde{x}(\varepsilon))$ is concave. Hence, we can apply the pipage rounding directly: We start with the fractional solution x^*. At every step, we pick e' and e'' (using the algorithm from [7]) and move to $\tilde{x}(\varepsilon)$ with $\varepsilon = \varepsilon_1 = -\min\{x_{e'}, 1 - x_{e''}\}$ or $\varepsilon = \varepsilon_2 = \min\{1 - x_{e'}, x_{e''}\}$ whichever minimizes the concave function $F(\tilde{x}(\varepsilon))$ on the interval $[\varepsilon_1, \varepsilon_2]$. We stop when the current solution $\tilde{x}$ is integral.

At every step, we decrease the number of fractional variables x_e by at least 1. Thus, we terminate the algorithm in at most $|E|$ iterations. The value of the function $F(\tilde{x})$ never increases. So the cost of the final integral solution is at most the cost of the initial fractional solution x^*, which, in turn, is at most $A_q \cdot LP^*$.

Note, that we have not used any special properties of graphic matroids. The algorithm from [7] works for general matroids accessible through oracle calls. So we can apply our technique to more

general problems where the objective is to minimize a function like (21) subject to base matroid constraints.

3 Proof of Theorem 1.4

We now give an efficient algorithm for finding $(1+\varepsilon)$ approximately optimal solution to LP (6)-(10).

Proof of Theorem 1.4. Observe that for every $y \in \mathcal{P}$, there exists a z such that the pair (y, z) is a feasible solution to LP (6)-(10). For example, one such z is defined as $z_{jS} = \prod_{i \in S} y_i \prod_{i \notin S}(1 - y_i)$. Of course, this particular z may be suboptimal. However, it turns out, as we show below, that for every y, we can find the optimal z efficiently. Let us denote the minimal cost of the j-th term in (6) for a given $y \in \mathcal{P}$ by $H_j(y)$. That is, $H_j(y)$ is the cost of the following LP. The variables of the LP are z_{jS}. The parameters $y \in \mathcal{P}$ and $j \in [k]$ are fixed.

$$\min \sum_{S \subseteq [n]} f_j\Big(\sum_{i \in S} d_{ij}\Big) z_{jS} \tag{22}$$

$$\sum_{S \subseteq [n]} z_{jS} = 1 \tag{23}$$

$$\sum_{S : i \in S} z_{jS} = y_i, \qquad \forall i \in [n] \tag{24}$$

$$z_{jS} \geq 0, \qquad \forall S \subseteq [n] \tag{25}$$

Now, LP (6)-(10) can be equivalently rewritten as (below y is the variable).

$$\min \sum_{j \in [k]} H_j(y) \tag{26}$$

$$y \in \mathcal{P} \tag{27}$$

The functions $H_j(y)$ are convex[2]. In Lemma 3.1 (see below), we prove that LP (22)-(25) can be solved in polynomial time, and thus the functions $H_j(y)$ can be computed efficiently. The algorithm for finding $H_j(y)$ also returns a subgradient of H_j at y. Hence, the minimum of convex problem (26)-(27) can be found using the ellipsoid method. Once the optimal y^* is found, we find z^* by solving LP (22)-(25) for y^* and each $j \in [k]$. □

Lemma 3.1. *There exists a polynomial time algorithm for computing H_j and finding a subgradient of H_j.*

Proof. We need to solve LP (22)-(25). Recall that in Section D (Theorem D.1) we show how to choose $\delta_j > 0$ such that each d_{ij}/δ_j is an integer polynomially bounded. For simplicity we assume that $\delta_j = 1$ (the proof in general case is almost identical). Therefore d_{ij} are integral in this case and polynomially bounded. We write the dual LP. We introduce a variable ξ for constraint (23) and variables η_i for constraints (24).

$$\max \;\; \xi + \sum_i \eta_i y_i \tag{28}$$

$$\xi + \sum_{i \in S} \eta_i \leq f_j\Big(\sum_{i \in S} d_{ij}\Big), \qquad \forall S \subset [n] \tag{29}$$

[2]If z^* and z^{**} are the optimal solutions for vectors y^* and y^{**}, then $\lambda z^* + (1-\lambda) z^{**}$ is a feasible solution for $\lambda y^* + (1-\lambda) y^{**}$. Hence, $H_j(\lambda y^* + (1-\lambda) y^{**}) \leq \lambda H_j(y^*) + (1-\lambda) H_j(y^{**})$. See Section C in Appendix for details.

The LP has exponentially many constraints. However, finding a violated constraint is easy. To do so, we guess $B^* = \sum_{i \in S^*} d_{ij}$ for the set S^* violating the constraint. That is possible, since all d_{ij} are polynomially bounded, and so is B^*. Then we solve the maximum knapsack problem

$$\max_{S \subseteq [n]} \sum_{i \in S} \eta_i$$
$$\sum_{i \in S} d_{ij} = B^*$$

using the standard dynamic programming algorithm and obtain the optimal set S^*. The knapsack problem is polynomially solvable, since B^* is polynomially bounded. If $\xi + \sum_{i \in S^*} \eta_i > f_j(B^*)$, then constraint (29) is violated for the set S^*; otherwise all constraints (29) are satisfied.

Let (ξ^*, η^*) be the optimal solution of the dual LP. The value of the function $H_j(y)$ equals the objective value of the dual LP. A subgradient of H_j at y is given by the equation

$$\tilde{y} \mapsto \xi^* + \sum_i \eta_i^* \tilde{y}_i. \tag{30}$$

This is a subgradient of H_j, since (ξ^*, η^*) is a feasible solution of the dual LP for every $\tilde{y}$ (note that constraint (29) does not depend on y), and, hence, (30) is a lower bound on $H_j(\tilde{y})$. □

4 Proof of Theorem 1.5

In this section, we prove the main theorem – Theorem 1.5.

Proof of Theorem 1.5. The theorem easily follows from the de la Peña decoupling inequality (Theorem 1.2 and Corollary 6.2) for $f_j(t) = c_j t^q$, and from the more general inequality presented in Theorem 5.3 (see also Corollary 5.4) for arbitrary convex functions f_j. Consider a feasible solution (y^*, z^*) to IP (3)-(5). We prove inequality (12) term by term. That is, for every j we show that

$$\mathbb{E}\Big[f_j\Big(\sum_{i \in D_j} d_{ij} R_i(y^*)\Big)\Big] \le A(f_j) \sum_{S \subseteq [n]} f_j\Big(\sum_{i \in S \cap D_j} d_{ij}\Big) z^*_{jS}. \tag{31}$$

Recall that $D_j = \{i : d_{ij} \neq 0\}$. Above, we dropped terms with $i \notin D_j$, since if $i \notin D_j$, then $d_{ij} = 0$.

Fix a $j \in [n]$. Define random variables Y_i for $i \in D_j$ as follows: Pick a random set $S \subset [n]$ with probability z_{jS}, and let $Y_i = d_{ij}$ if $i \in S$, and $Y_i = 0$ otherwise. Note that random variables Y_i are dependent. We have

$$\Pr(Y_i = d_{ij}) = \sum_{S : i \in S} \Pr(S) = \sum_{S : i \in S} z^*_{jS} = y^*_i.$$

It is easy to see that

$$\mathbb{E}\Big[f_j\Big(\sum_{i \in [n]} Y_i\Big)\Big] = \sum_{S \subseteq [n]} f_j\Big(\sum_{i \in S} d_{ij}\Big) z_{jS}.$$

The right hand side is simply the definition of the expectation on the left hand side. Now, let $X_i = d_{ij}\, R_i(y^*)$ for $i \in D_j$. Note that by conditions of the theorem, $\Pr(X_i = d_{ij}) = y^*_i = \Pr(Y_i = d_{ij})$ (by condition (1)). Thus, each X_i has the same distribution as Y_i. Furthermore, X_i's are independent or negatively associated (by condition (2)). Therefore, we can apply the decoupling inequality from Theorem 5.3 (see Corollary 5.4)

$$\mathbb{E}\Big[f_j\Big(\sum_{i \in D_j} X_i\Big)\Big] \le A(f_j) \mathbb{E}\Big[f_j\Big(\sum_{i \in D_j} Y_i\Big)\Big].$$

The left hand side of the inequality equals the left hand side of (31), the right hand side of the inequality equals the right hand side of (31). Hence, inequality (31) holds. □

5 Decoupling Inequality

In this section, we prove the decoupling inequality (Theorem 1.2) with the optimal constant $C_q = A_q^{1/q}$. In fact, we prove a more general inequality which works for arbitrary convex functions. To state the inequality we need the notion of *convex stochastic order*.

Definition 5.1. *We say that a random variable X is less than Y in the convex (stochastic) order, and write $X \leq_{cx} Y$ if for every convex function $\varphi : \mathbb{R} \to \mathbb{R}$,*

$$\mathbb{E}\varphi(X) \leq \mathbb{E}\varphi(Y), \tag{32}$$

whenever both expectations exist.

Remark 5.2. *If $X \leq_{cx} Y$ and ψ is a* concave *function, then $\mathbb{E}[\psi(X)] \geq \mathbb{E}[\psi(Y)]$, since the function $\phi(x) = -\psi(x)$ is convex, and therefore $-\mathbb{E}[\psi(X)] \leq -\mathbb{E}[\psi(Y)]$.*

It is easy to see that the convex stochastic order defines a partial order on all random variables. Particularly, if $X \leq_{cx} Y$ and $Y \leq_{cx} Z$, then $X \leq_{cx} Z$. Note that the definition depends only on the distributions of X and Y. That is, if X has the same distribution as Z and $X \leq_{cx} Y$, then $Z \leq_{cx} Y$. The random variables X and Y may be defined on the same probability space or on different probability spaces. We refer the reader to the book of Shaked and Shanthikumar [32] for a detailed introduction to stochastic orders.

We now state the general inequality in terms of the convex order. Part II of the theorem shows that the inequality is tight.

Theorem 5.3. *I. Let $Y_1, \ldots, Y_n$ be jointly distributed nonnegative (non-independent) random variables, and let $X_1, \ldots, X_n$ be independent random variables such that each X_i has the same distribution as Y_i. Let P be a Poisson random variable with parameter 1 independent of Y_i's. Then,*

$$\sum_{i=1}^{n} X_i \leq_{cx} P \sum_{i=1}^{n} Y_i. \tag{33}$$

II. For every nonnegative function φ with a finite expectation $\mathbb{E}[\varphi(P)]$ and every positive ε, there exists n and random variables $X_1, \ldots, X_n$, $Y_1, \ldots, Y_n$ as in part I such that

$$\mathbb{E}\left[\varphi\Big(\sum_{i=1}^{n} X_i\Big)\right] \geq (1+\varepsilon)\mathbb{E}\left[\varphi\Big(P\sum_{i=1}^{n} Y_i\Big)\right]. \tag{34}$$

Proof of Theorem 1.2. Theorem 5.3 implies Theorem 1.2, since the function $t \to t^q$ is convex and thus

$$\begin{aligned}
\|X_1 + \cdots + X_n\|_q^q &= \mathbb{E}\left[(X_1 + \cdots + X_n)^q\right] \leq \mathbb{E}\left[\big(P \cdot (Y_1 + \cdots + Y_n)\big)^q\right] \\
&= \mathbb{E}\left[P^q\right] \cdot \mathbb{E}\left[(Y_1 + \cdots + Y_n)^q\right] = A_q \mathbb{E}\left[(Y_1 + \cdots + Y_n)^q\right] \\
&= A_q \|Y_1 + \cdots + Y_n\|_q^q.
\end{aligned}$$

Part II of Theorem 5.3 shows that we cannot replace A_q with a smaller constant. □

Another immediate corollary of Theorem 5.3 is as follows.

Corollary 5.4. *For an arbitrary nonnegative convex function f,*

$$\mathbb{E}[f(X_1+\cdots+X_n)] \le A(f)\mathbb{E}[f(Y_1+\cdots+Y_n)],$$

where $A(f)$ is defined in (11).

Proof. Write,

$$\mathbb{E}[f(X_1+\cdots+X_n)] \le \mathbb{E}\big[f\big(P\cdot(Y_1+\cdots+Y_n)\big)\big] = \mathbb{E}\mathbb{E}\big[f\big(P\cdot(Y_1+\cdots+Y_n)\big) \mid Y_1+\cdots+Y_n\big].$$

For any $t>0$, particularly for $t=\sum_i Y_i$, we have $\mathbb{E}[f(tP)] \le A(f)f(t)$, hence

$$\mathbb{E}[f(X_1+\cdots+X_n)] \le \mathbb{E}[A(f)f(Y_1+\cdots+Y_n)] = A(f)\mathbb{E}[f(Y_1+\cdots+Y_n)].$$

□

We first prove part II of Theorem 5.3. Consider the following example. Let $Y_i^{(n)}, i\in\{1,\dots,n\}$, be random variables taking value 1 with probability $1/n$, and 0 with probability $1-1/n$. We generate $Y_i^{(n)}$'s as follows. We pick a random $j\in[n]$ and let $Y_j^{(n)}=1$ and $Y_i^{(n)}=0$ for $i\neq j$. Random variables $X_i^{(n)}$ are i.i.d. Bernoulli random variables with $\mathbb{E}[X_i^{(n)}]=1/n$. Then, the sum $\sum_{i=1}^n Y_i^{(n)}$ always equals 1, and $\mathbb{E}\left[\varphi(P\sum_{i=1}^n Y_i^{(n)})\right] = \mathbb{E}\left[\varphi(P)\right]$. As $n\to\infty$, the sum $\sum_{i=1}^n X_i^{(n)}$ converges in distribution to P (by the Poisson limit theorem). Thus (see Lemma E.1 for details),

$$\sup_n \mathbb{E}\left[\varphi\Big(\sum_{i=1}^n X_i^{(n)}\Big)\right] \ge \mathbb{E}\left[\varphi(P)\right],$$

and, hence, for some n inequality (34) holds.

Before proceeding to the proof of part I, we state some known properties of the convex order.

Lemma 5.5 (Theorem 3.A.12 in [32])**.** *Suppose $X_1,\dots,X_n$ are independent random variables and $Y_1,\dots,Y_n$ are independent random variables. If $X_i \le_{cx} Y_i$ for all i, then $\sum_{i=1}^n X_i \le_{cx} \sum_{i=1}^n Y_i$.*

Lemma 5.6 (Theorem 3.A.36 in [32])**.** *Consider random variables $X_1,\dots,X_n$ and Y. If $X_i \le_{cx} Y$ for all i, then $\sum_{i=1}^n a_i X_i \le \sum_{i=1}^n a_i Y$, for any sequence of nonnegative numbers $a_1,\dots,a_n$.*

For completeness, we prove Lemma 5.5 and Lemma 5.6 in Appendix. To simplify the proof, we will use the following easy lemma.

Lemma 5.7. *If for random variables X and Y condition (32) is satisfied for all convex functions φ with $\varphi(0)=0$, then $X\le_{cx} Y$.*

Proof. Consider an arbitrary convex function φ. Let $\tilde\varphi(x)=\varphi(x)-\varphi(0)$. Since $\tilde\varphi(0)=0$, we have

$$\mathbb{E}[\varphi(X)] = \mathbb{E}[\tilde\varphi(X)]+\varphi(0) \le \mathbb{E}[\tilde\varphi(Y)]+\varphi(0) = \mathbb{E}[\varphi(Y)].$$

□

First, we prove that a Bernoulli random variable B can be upper bounded in the convex order by a Poisson random variable P with $\mathbb{E}[P]=\mathbb{E}[B]$.

Lemma 5.8. *Let P be an integral random variable, and B be a Bernoulli random variable with parameter $p = \mathbb{E}[P]$. Then, $B \leq_{cx} P$. Particularly, if P is a Poisson random variable with $\mathbb{E}[P] = \mathbb{E}[B]$, then $B \leq_{cx} P$.*

Proof. Consider an arbitrary convex function $\varphi : \mathbb{R} \to \mathbb{R}$ with $\varphi(0) = 0$ (see Lemma 5.7). Define linear function $l : \mathbb{R} \to \mathbb{R}$ as $l(x) = \varphi(1)x$. The graph of l intersects the graph of φ at points $(0, 0)$ and $(1, \varphi(1))$. Since φ is convex, we have $l(x) \leq \varphi(x)$ for $x \notin (0, 1)$. Hence, for every integral k, $l(k) \leq \varphi(k)$, and $l(P) \leq \varphi(P)$. Consequently,

$$\mathbb{E}\left[\varphi(P)\right] \geq \mathbb{E}\left[l(P)\right] = l(\mathbb{E}\left[P\right]) = l(\mathbb{E}\left[B\right]) = \varphi(1)\mathbb{E}\left[B\right] = \mathbb{E}[\varphi(B)].$$

□

Now we consider a very special case of Theorem 5.3 when all X_i's and Y_i's are Bernoulli random variables scaled by a factor α_i, and all events $\{Y_i = 1\}$ are mutually exclusive. As we see later, the general case can be easily reduced to this special case.

Lemma 5.9. *Consider Bernoulli random variables $\chi_1, \ldots, \chi_n$ such that $\sum_{i=1}^{n} \Pr(\chi_i = 1) = 1$, and all events $\{\chi_i = 1\}$ are mutually exclusive. (That is, with probability 1 one and only one χ_i equals 1.) Let $B_1, \ldots, B_n$ be independent Bernoulli random variables such that $\Pr(B_i = 1) = \Pr(\chi_i = 1)$. Then, for all nonnegative numbers $\alpha_1, \ldots, \alpha_n$, we have*

$$\sum_{i=1}^{n} \alpha_i \chi_i \leq_{cx} \sum_{i=1}^{n} \alpha_i B_i \leq_{cx} P \sum_{i=1}^{n} \alpha_i \chi_i,$$

where P is a Poisson random variable with parameter 1 independent of χ_i's.

Proof. I. We prove the first inequality. Consider an arbitrary convex function $\varphi : \mathbb{R} \to \mathbb{R}$ with $\varphi(0) = 0$. The function φ is superadditive i.e., $\varphi(a + b) \geq \varphi(a) + \varphi(b)$ for all positive a and b (since the derivative of φ is monotonically non-decreasing, and $\varphi(0) = 0$). Hence, we have

$$\varphi\big(\sum_{i=1}^{n} \alpha_i B_i\big) \geq \sum_{i=1}^{n} \varphi(\alpha_i B_i) = \sum_{i=1}^{n} B_i \varphi(\alpha_i),$$

and

$$\mathbb{E}\left[\varphi\big(\sum_{i=1}^{n} \alpha_i B_i\big)\right] \geq \mathbb{E}\left[\sum_{i=1}^{n} B_i \varphi(\alpha_i)\right] = \sum_{i=1}^{n} \Pr(B_i = 1)\varphi(\alpha_i)$$
$$= \sum_{i=1}^{n} \Pr(\chi_i = 1)\varphi(\alpha_i) = \mathbb{E}\left[\varphi\big(\sum_{i=1}^{n} \alpha_i \chi_i\big)\right].$$

II. We prove the second inequality. Using Lemma 5.8 and Lemma 5.5, we replace Bernoulli random variable $B_1, \ldots, B_n$ with independent Poisson random variables $P_1, \ldots, P_n$ satisfying $\mathbb{E}\, P_i = \mathbb{E}\, B_i$. The sum $P_1 + \cdots + P_n$ is distributed as a Poisson random variable with parameter 1. By coupling random variables $(P_1 + \cdots + P_n)$ and P, we may assume that $P = P_1 + \cdots + P_n$.

Consider an arbitrary convex function φ with $\varphi(0) = 0$. Using convexity of the function φ, we derive

$$\begin{aligned}
\mathbb{E}\left[\varphi\big(\sum_{i=1}^{n}\alpha_i P_i\big)\right] &= \sum_{k=1}^{\infty}\mathbb{E}\left[\varphi\big(\sum_{i=1}^{n}\alpha_i P_i\big)\ \middle|\ P=k\right]\Pr(P=k)\\
&= \sum_{k=1}^{\infty}\mathbb{E}\left[\varphi\big(\sum_{i=1}^{n}\frac{P_i}{k}\cdot\alpha_i k\big)\ \middle|\ P=k\right]\Pr(P=k)\\
&\le \sum_{k=1}^{\infty}\mathbb{E}\left[\sum_{i=1}^{n}\frac{P_i}{k}\varphi(\alpha_i k)\ \middle|\ P=k\right]\Pr(P=k)\\
&= \sum_{k=1}^{\infty}\sum_{i=1}^{n}\mathbb{E}\left[\frac{P_i}{k}\ \middle|\ P=k\right]\varphi(\alpha_i k)\Pr(P=k).
\end{aligned}$$

We observe that $\mathbb{E}\left[P_i \mid P=k\right] = k\,\mathbb{E}\left[P_i\right]$, which follows from the following well known fact (see e.g., Feller [14], Section IX.9, Problem 6(b), p. 237).

Fact 5.10. *Suppose P_a and P_b are independent Poisson random variables with parameters λ_a and λ_b. Then, for every $k \in \mathbb{N}$, $\mathbb{E}\big[P_a \mid P_a + P_b = k\big] = \frac{\lambda_a}{\lambda_a+\lambda_b}\,k$.*

In our case, $P_a = P_i$, $P_b = \sum_{i'\neq i} P_{i'}$, $P_a + P_b = P$. Therefore, we have

$$\begin{aligned}
\mathbb{E}\Big[\varphi\big(\sum_{i=1}^{n}\alpha_i P_i\big)\Big] &\le \sum_{k=1}^{\infty}\sum_{i=1}^{n}\mathbb{E}\Big[\frac{P_i}{k}\ \Big|\ P=k\Big]\varphi(\alpha_i k)\Pr(P=k)\\
&= \sum_{k=1}^{\infty}\sum_{i=1}^{n}\mathbb{E}\left[P_i\right]\varphi(\alpha_i k)\Pr(P=k)\\
&= \sum_{i=1}^{n}\mathbb{E}\left[\chi_i\right]\sum_{k=1}^{\infty}\varphi(\alpha_i k)\Pr(P=k)\\
&= \sum_{i=1}^{n}\mathbb{E}\left[\chi_i\right]\mathbb{E}\left[\varphi(\alpha_i k)\right]\\
&= \mathbb{E}\Big[\varphi(P\sum_{i=1}^{n}\alpha_i\chi_i)\Big].
\end{aligned}$$

□

Proof of Theorem 5.3. Denote by $\mathcal{Y}$ the support of the random vector $Y = (Y_1, \ldots, Y_n)$. Each Y_i can be represented as follows

$$Y_i = \sum_{y\in\mathcal{Y}} y_i\,\chi(Y=y),$$

where $\chi(Y=y)$ is the indicator of the event $\{Y=y\}$. Here we assume that $\mathcal{Y}$ is finite. We treat the general case in Appendix G. Applying Lemma 5.9 to the random variables $\chi(Y=y)$, we get

$$Y_i = \sum_{y\in\mathcal{Y}} y_i\,\chi(Y=y) \le_{cx} \sum_{y\in\mathcal{Y}} y_i\,B^i_y,$$

where B_y^i are independent Bernoulli random variables. Each B_y^i is distributed as the random variable $\chi(Y = y)$; that is, $\Pr(B_y^i = 1) = \Pr(Y = y)$. Since each X_i has the same distribution as Y_i, we have

$$X_i \leq_{cx} \sum_{y \in \mathcal{Y}} y_i \, B_y^i.$$

By Lemma 5.5, we can sum up this inequality over all i from 1 to n:

$$\sum_{i=1}^{n} X_i \leq_{cx} \sum_{i=1}^{n} \sum_{y \in \mathcal{Y}} y_i \, B_y^i = \sum_{y \in \mathcal{Y}} \Big(\sum_{i=1}^{n} y_i \, B_y^i \Big).$$

We apply Lemma 5.6 to every sum in parentheses:

$$\sum_{i=1}^{n} y_i \, B_y^i \leq_{cx} \sum_{i=1}^{n} y_i \, B_y.$$

Again, using Lemma 5.5, we get

$$\sum_{i=1}^{n} X_i \leq_{cx} \sum_{y \in \mathcal{Y}} \Big(\sum_{i=1}^{n} y_i \, B_y^i \Big) \leq_{cx} \sum_{y \in \mathcal{Y}} \Big(\sum_{i=1}^{n} y_i \Big) \, B_y.$$

Finally, by Lemma 5.9,

$$\sum_{y \in \mathcal{Y}} \Big(\sum_{i=1}^{n} y_i \Big) \, B_y \leq_{cx} P \cdot \sum_{y \in \mathcal{Y}} \Big(\sum_{i=1}^{n} y_i \Big) \, \chi(Y = y) = P \cdot \sum_{i=1}^{n} \Big(\sum_{y \in \mathcal{Y}} y_i \, \chi(Y = y) \Big) = P \sum_{i=1}^{n} Y_i.$$

This concludes the proof of Theorem 5.3. □

6 Negatively Associated Random Variables

The decoupling inequalities (2) and (33) can be extended to *negatively associated* random variables $X_1, \ldots, X_n$. The notion of negative association is defined as follows.

Definition 6.1 (Joag-Dev and Proschan [21])**.** *Random variables $X_1, \ldots, X_n$ are negatively associated if for all disjoint sets $I, J \subset [n]$ and all non-decreasing functions $f : \mathbb{R}^I \to \mathbb{R}$ and $g : \mathbb{R}^J \to \mathbb{R}$ the following inequality holds:*

$$\mathbb{E}\left[f(X_i, i \in I) \cdot g(X_j, j \in J) \right] \leq \mathbb{E}\left[f(X_i, i \in I) \right] \cdot \mathbb{E}\left[g(X_j, j \in J) \right].$$

Shao [33] showed that if $X_1, \ldots, X_n$ are negatively associated random variables, and $X_1^*, \ldots, X_n^*$ are independent random variables such that each X_i^* is distributed as X_i, then for every convex function $\varphi : \mathbb{R} \to \mathbb{R}$,

$$\mathbb{E}\left[\varphi(X_1 + \cdots + X_n) \right] \leq \mathbb{E}\left[\varphi(X_1^* + \cdots + X_n^*) \right].$$

In other words, $X_1 + \cdots + X_n \leq_{cx} X_1^* + \cdots + X_n^*$ (see also Theorem 3.A.39 in [32]).

Corollary 6.2. *Let $Y_1, \ldots, Y_n$ be jointly distributed nonnegative (non-independent) random variables, and let $X_1, \ldots, X_n$ be negatively associated random variables such that each X_i has the same distribution as Y_i. Let P be a Poisson random variable with parameter 1 independent of the random variables Y_i. Then,*

$$\sum_{i=1}^{n} X_i \leq_{cx} P \sum_{i=1}^{n} Y_i.$$

Particularly, for every convex nonnegative f,

$$f\big(\sum_{i=1}^{n} X_i\big) \le A(f) f\big(\sum_{i=1}^{n} Y_i\big),$$

and for $q \ge 1$,

$$\big\|\sum_{i=1}^{n} X_i\big\|_q \le A_q^{1/q} \big\|\sum_{i=1}^{n} Y_i\big\|_q,$$

where $A(f)$ is defined in (11), and A_q is the fractional Bell number.

7 Generalizations

We can extend our results to maximization problems with the objective function

$$\sum_{j\in[k]} f_j\big(\sum_{i\in[n]} d_{ij} y_i\big), \tag{35}$$

if f_j's are arbitrary non-decreasing nonnegative *concave* functions defined on $\mathbb{R}^{\ge 0}$. The approximation ratio equals $\min_j B(f_j)$, where

$$B(f_j) = \inf_{t>0} \mathbb{E}\Big[\frac{f_j(Pt)}{f_j(t)}\Big].$$

It is not hard to see that $B(f) \ge 1 - 1/e$ for all f. Indeed, if $P \ge 1$, then $f(Pt)/f(t) \ge 1$, thus $B(f) = \mathbb{E}[f(Pt)/f(t)] \ge \Pr(P \ge 1) = 1 - 1/e$. This bound is tight if $f(Pt)/f(t) = 1$ for $P \ge 1$. For example, $B(f) = 1 - 1/e$ for the function $f(t) = \min\{t, 1\}$. Note that the approximation ratio of $1 - 1/e \approx 0.632$ for maximization problems of this form was previously known (see Calinescu et al. [7]). However, for some concave functions f we get a better approximation. For example, for $f(t) = \sqrt{t}$, we get an approximation ratio of $B(\sqrt{t}) \approx 0.773$.

Acknowledgment

We would like to thank the anonymous referees for valuable comments.

References

[1] Alexander A. Ageev and Maxim Sviridenko. Pipage Rounding: A New Method of Constructing Algorithms with Proven Performance Guarantee. J. Comb. Optim. 8(3): 307-328, (2004).

[2] S. Albers. Energy-efficient algorithms. Commun. ACM 53(5), 86-96 (2010).

[3] Matthew Andrews, Antonio Fernandez Anta, Lisa Zhang, Wenbo Zhao. Routing for Power Minimization in the Speed Scaling Model. IEEE/ACM Trans. Netw. 20(1): 285-294 (2012).

[4] Matthew Andrews, Spyridon Antonakopoulos, Lisa Zhang. Minimum-Cost Network Design with (Dis)economies of Scale. FOCS 2010, pp. 585-592.

[5] A. Assad and W. Xu. The quadratic minimum spanning tree problem. Naval Research Logistics v39. (1992), pp. 399-417.

[6] Yossi Azar, Amir Epstein. Convex programming for scheduling unrelated parallel machines. STOC 2005, pp. 331-337.

[7] Gruia Calinescu, Chandra Chekuri, Martin Pal and Jan Vondrak. Maximizing a submodular set function subject to a matroid constraint. SIAM Journal on Computing 40:6 (2011), pp. 1740-1766.

[8] V. Chvatal. On certain polytopes associated with graphs. J. Combinatorial Theory Ser. B, 18 (1975), pp. 138-154.

[9] Evripidis Bampis, Alexander Kononov, Dimitrios Letsios, Giorgio Lucarelli and Maxim Sviridenko. Energy Efficient Scheduling and Routing via Randomized Rounding. FSTTCS 2013, pp. 449-460.

[10] Nikhil Bansal and Kirk Pruhs. The Geometry of Scheduling. FOCS 2010, pp. 407-414.

[11] Stephen Boyd and Lieven Vandenberghe. Convex Optimization. Cambridge University Press, 2004.

[12] C. Durr and O. Vasquez. Order constraints for single machine scheduling with non-linear cost. The 16th Workshop on Algorithm Engineering and Experiments (ALENEX), 2014.

[13] M. Grotschel, L. Lovasz and A. Schrijver. Geometric algorithms and combinatorial optimization. Springer-Verlag, Berlin, 1988.

[14] W. Feller. An Introduction to Probability Theory and its Applications (3rd edition). John Wiley & Sons, New York (1968).

[15] D. R. Fulkerson. Blocking and anti-blocking pairs of polyhedra. Math. Programming, 1 (1971), pp. 168-194.

[16] Michel X. Goemans. Minimum Bounded Degree Spanning Trees. FOCS 2006, 273-282.

[17] W. Hohn and T. Jacobs. An experimental and analytical study of order constraints for single machine scheduling with quadratic cost. 14th Workshop on Algorithm Engineering and Experiments (ALENEX), pp. 103-117, SIAM, 2012.

[18] W. Hohn and T. Jacobs. On the performance of Smith's rule in single-machine scheduling with nonlinear cost. LATIN 2012, pp. 482-493.

[19] Han Hoogeveen, Petra Schuurman, Gerhard J. Woeginger. Non-Approximability Results for Scheduling Problems with Minsum Criteria. INFORMS Journal on Computing 13(2), pp. 157-168 (2001).

[20] S. Irani and K. Pruhs. Algorithmic problems in power management. ACM SIGACT News 36 (2), 63-76.

[21] Kumar Joag-Dev and Frank Proschan. Negative Association of Random Variables with Applications. The Annals of Statistics 11 (1983), no. 1, 286-295.

[22] V. S. A. Kumar, M. V. Marathe, S. Parthasarathy and A. Srinivasan. A Unified Approach to Scheduling on Unrelated Parallel Machines. Journal of the ACM, Vol. 56, 2009.

[23] S. Maia, E. Goldbarg, and M. Goldbarg. On the biobjective adjacent only quadratic spanning tree problem. Electronic Notes in Discrete Mathematics, 41, (2013), 535-542.

[24] Yurii Nesterov and Arkadi Nemirovski. Interior-point polynomial algorithms in convex programming. Society for Industrial and Applied Mathematics, 1987.

[25] T. Oncan and A. Punnen. The quadratic minimum spanning tree problem: A lower bounding procedure and an efficient search algorithm. Computers and Operations Research v.37, (2010), 1762-1773.

[26] Victor de la Peña. Bounds on the expectation of functions of martingales and sums of positive RVs in terms of norms of sums of independent random variables. Proceedings of the American Mathematical Society 108.1 (1990): 233-239.

[27] Victor de la Peña and E. Giné. Decoupling: from dependence to independence. Springer, 1999.

[28] Victor de la Peña, Rustam Ibragimov, Shaturgan Sharakhmetov. On Extremal Distributions and Sharp L_p-Bounds for Sums of Multilinear Forms. The Annals of Probability, 2003, vol. 31 (2),pp. 630–675.

[29] D. Pereira, M. Gendreau, A. Cunha. Stronger lower bounds for the quadratic minimum spanning tree problem with adjacency costs. Electronic Notes in Discrete Mathematics, 41, (2013), 229-236.

[30] A. Schrijver. Theory of Linear and Integer Programming. Wiley, 1998.

[31] A. Schrijver. Combinatorial Optimization. Springer, 2002.

[32] Moshe Shaked and J. George Shanthikumar. Stochastic orders. Springer, 2007.

[33] Qi-Man Shao. A comparison theorem on moment inequalities between negatively associated and independent random variables. Journal of Theoretical Probability 13, no. 2 (2000): 343-356.

[34] Andreas S. Schulz and Martin Skutella. Scheduling Unrelated Machines by Randomized Rounding. SIAM J. Discrete Math, 15(4), pp. 450-469 (2002).

[35] Mohit Singh and Lap Chi Lau. Approximating minimum bounded degree spanning trees to within one of optimal. STOC 2007: 661-670.

[36] Martin Skutella. Convex quadratic and semidefinite programming relaxations in scheduling. J. ACM 48(2), 206-242 (2001).

[37] Martin Skutella. Approximation and randomization in scheduling. Technische Universitat Berlin, Germany, 1998.

[38] S. Stiller and A. Wiese. Increasing Speed Scheduling and Flow Scheduling. ISAAC (2) 2010, pp.279-290.

[39] A. Wierman, L. L. H. Andrew, and A. Tang. Power-aware speed scaling in processor sharing systems. INFOCOM 2009, pp. 2007-2015.

A Corollary A.1

In this section, we prove a corollary of Theorem 5.3, which we will need in the next section.

Corollary A.1. *Let $Y_1, \dots, Y_n$ be jointly distributed (non-independent) nonnegative integral random variables, and let $X_1, \dots, X_n$ be independent Bernoulli random variables taking values 0 and 1 such that for each i, $\mathbb{E}[X_i] = \mathbb{E}[Y_i]$. Then,*

$$X_1 + \cdots + X_n \leq_{cx} P(Y_1 + \cdots + Y_n).$$

Particularly, for every $q \geq 1$,

$$\| \sum_{i=1}^{n} X_i \|_q \leq A_q^{1/q} \| \sum_{i=1}^{n} Y_i \|_q,$$

where A_q is the fractional Bell number.

Proof. Consider independent random variables X_i^* such that each X_i^* is distributed as Y_i. By Lemma 5.8, $X_i \leq_{cx} X_i^*$ for all i. Hence, by Lemma 5.5 and Theorem 5.3,

$$X_1 + \cdots + X_n \leq_{cx} X_1^* + \cdots + X_n^* \leq_{cx} Y_1 + \cdots + Y_n.$$

□

B Unrelated Parallel Machine Scheduling with Nonlinear Functions of Completion Times – Technical Details

We consider the following linear programming relaxation of the scheduling problem $R||\sum_j w_j C_j^p$. The variable $x_{ijt} = 1$ if job j starts at time t on machine i. For convenience we assume that $x_{ijt} = 0$ for negative values of t.

$$\min \sum_{i\in[m]} \sum_{j\in[n]} \sum_{t\geq 0} w_j (t + p_{ij})^p x_{ijt} \tag{36}$$

$$\sum_{i\in[m], t\geq 0} x_{ijt} = 1 \qquad \forall j \in [n] \tag{37}$$

$$\sum_{j\in[n]} \sum_{\tau=t-p_{ij}+1}^{t} x_{ij\tau} \leq 1 \qquad \forall i \in [m], t \geq 0 \tag{38}$$

$$x_{ijt} \geq 0 \qquad \forall i \in [m], j \in [n], t \geq 0. \tag{39}$$

The constraints (37) say that each job must be assigned, the constraint (38) says that at most one job can be processed in a unit time interval on each machine. Such linear programming relaxation are known under the name of *strong time indexed formulations*. The standard issue with such relaxations is that they have pseudo-polynomially many variables due to potentially large number of indices t. One way to handle this issue is to partition the time interval into intervals $((1+\varepsilon)^k, (1+\varepsilon)^{k+1}]$ and round all completion times to the endpoints of such intervals. This method leads to polynomially sized linear programming relaxations with $(1 + O(\varepsilon))$-loss in the performance guarantee (see [37] for detailed description of the method). From now on we ignore this issue and assume that the planning horizon upper bound $\sum_{i,j} p_{ij}$ is polynomially bounded in the input size.

Algorithm. Our approximation algorithm solves linear programming relaxation (36)-(39). Let x^* be the optimal fractional solution of the LP. Each job is tentatively assigned to machine i to

start at time t with probability x^*_{ijt}, independently at random. Let t_j be the tentative start time assigned to job j by our randomized procedure. We process jobs assigned to each machine in the order of the tentative completion times $t_j + p_{ij}$.

Analysis. We estimate the expected cost of the approximate solution returned by the algorithm. We denote the expected cost by APX. For each machine-job-tentative time triple (i, j, t), let J_{ijt} be the set of triples (i, j', t') such that $t' + p_{ij'} \le t + p_{ij}$. Let X^{ijt} be the random boolean variable such that $X^{ijt} = 1$ if job j is assigned to machine i with tentative start time t. In addition, let $Z^{ijt}_{j'}$ be the random boolean variable such that $Z^{ijt}_{j'} = 1$ if job j' is assigned to machine i with tentative start time t' for some $(i, j', t') \in J_{ijt}$ by our randomized rounding procedure. Then,

$$\Pr\left(Z^{ijt}_{j'} = 1\right) = \sum_{t':(i,j',t')\in J_{ijt}} x^*_{ij't'}.$$

Suppose that job j is tentative scheduled on machine i at time t i.e., $X^{ijt} = 1$. We start processing job j after all jobs j' tentative scheduled on machine i at time t' with $t' + p_{ij'} \le t + p_{ij}$ are finished. Thus the weighted expected completion time to the power of p for j equals (given $X^{ijt} = 1$)

$$\begin{aligned}\mathbb{E}\left[w_j C_j^p \;\middle|\; X^{ijt} = 1\right] &= \mathbb{E}\Big[\Big(\sum_{j'\in[n]\setminus\{j\}} p_{i,j'} Z^{ijt}_{j'} + p_{ij}\Big)^p \;\middle|\; X^{ijt} = 1\Big] \\ &= \mathbb{E}\Big[\Big(\sum_{j'\in[n]\setminus\{j\}} p_{i,j'} Z^{ijt}_{j'} + p_{ij}\Big)^p\Big].\end{aligned}$$

In the second equality, we used that random variables $Z^{ijt}_{j'}$ are independent from the random variable X^{ijt}. Then,

$$\begin{aligned}APX &= \sum_{j\in[n]}\sum_{i\in[m]}\sum_{t\ge 0} w_j \mathbb{E}\left[w_j C_j^p \;\middle|\; X^{ijt} = 1\right] \Pr\left(X^{ijt} = 1\right) \\ &= \sum_{j\in[n]}\sum_{i\in[m]}\sum_{t\ge 0} w_j \mathbb{E}\Big[\Big(\sum_{j'\in[n]\setminus\{j\}} p_{i,j'} Z^{ijt}_{ij'} + p_{ij}\Big)^p\Big]\, x^*_{ijt}. \qquad (40)\end{aligned}$$

Note, that for fixed $i \in [m], j \in [n], t \ge 0$ random variables $Z^{ijt}_{j'}$ are independent from each other. We claim that

$$\mathbb{E}\Big[\Big(\sum_{j'\in[n]\setminus\{j\}} p_{i,j'} Z^{ijt}_{j'} + p_{ij}\Big)^p\Big] \le A_p(t + 2p_{ij})^p. \qquad (41)$$

Combining (40) and (41), we derive that the performance guarantee of our approximation algorithm is at most $2^p A_p$. We now prove inequality (41).

Let G_{ijt} be the interval graph where the vertex set $V(G_{ijt})$ is the collection of intervals corresponding to triples in $(i, j', t') \in J_{ijt}$ such that $x^*_{ij't'} > 0$. More precisely, every triple $(i, j', t') \in J_{ijt}$ corresponds to the interval $I_{ij't'} = [t', t' + p_{ij'})$ with corresponding weight $x^*_{ij't'} > 0$. Let $\mathcal{I}$ be the collection of all independent sets in G_{ijt}. The interval graph G_{ijt} is perfect, and the weights $x^*_{ij't'}$ satisfy the constraints (38), so there is a collection of weights $\lambda_C \ge 0$, $C \in \mathcal{I}$ (for more formal argument see below) such that

$$\begin{aligned}\sum_{C\in\mathcal{I}} \lambda_C &= 1, \\ \sum_{C\in\mathcal{I}: I_{ij't'}\in C} \lambda_C &= x^*_{ij't'}, \qquad \forall (i, j', t') \in J_{ijt}\end{aligned}$$

Formally, the claim above follows from the polyhedral characterization of perfect graphs proved by Fulkerson [15] and Chvatal [8] (see also Schrijver's book [30], Section 9, Application 9.2 on p. 118) that a graph G is perfect if and only if its stable set polytope is defined by the system below:

$$\begin{aligned} \sum_{v \in C} x_v &\leq 1, \qquad \text{for each clique } C, \\ x(v) &\geq 0, \qquad \text{for each } v \in V. \end{aligned}$$

In the interval graph G_{ijt} all clique inequalities are included in the constraints (38) and therefore any set of weights $x^*_{ij't'}$ can be decomposed into a convex combination of independent sets in G_{ijt}.

We define a random variable $Y^{ijt}_{j'}$ as follows: Sample an independent set $C \in \mathcal{I}$ with probability λ_C and let

$$Y^{ijt}_{j'} = \left|\{I_{ij't'} \in C\}\right|.$$

Note that one job j' may have more than one interval $I_{ij't'}$ in the set C (for different t'). Random variables $Y^{ijt}_{j'}$ may be dependent but

$$\mathbb{E}[Y^{ijt}_{j'}] = \sum_{(i,j',t') \in J_{ijt}} x^*_{ij't'} = \mathbb{E}[Z^{ijt}_{j'}].$$

Therefore, by Corollary A.1 we have

$$\mathbb{E}\Big[\Big(\sum_{j' \in [n]\setminus\{j\}} p_{i,j'} Z^{ijt}_{j'} + p_{ij}\Big)^p\Big] \leq A_p \mathbb{E}\Big[\Big(\sum_{j' \in [n]\setminus\{j\}} p_{i,j'} Y^{ijt}_{j'} + p_{ij}\Big)^p\Big]. \tag{42}$$

Now, observe, that $\sum_{j' \in [n]\setminus\{j\}} p_{i,j'} Y^{ijt}_{j'}$ is always bounded by $t + p_{ij}$, because all intervals in C are disjoint (C is an independent set) and all intervals are subsets of $[0, t + p_{ij}]$. Hence,

$$\mathbb{E}\Big[\Big(\sum_{j' \in [n]\setminus\{j\}} p_{i,j'} Z^{ijt}_{j'} + p_{ij}\Big)^p\Big] \leq A_p \mathbb{E}[((t + p_{ij}) + p_{ij})^p] \leq A_p (t + 2p_{ij})^p,$$

which concludes the proof.

C Convexity of H_j

We show that functions H_j defined in Section 3 are convex.

Lemma C.1. *Fix real numbers $q \geq 1$ and $d_1, \ldots, d_n \geq 0$. Define a function $H : [0,1]^n \to \mathbb{R}^+$ as follows: $H(y)$ equals the optimal value of the following LP:*

$$\begin{aligned} \min \sum_{S \subseteq [n]} \big(\sum_{i \in S} d_i\big)^q z_S & \\ \sum_{S \subseteq [n]} z_S &= 1 \\ \sum_{S : i \in S} z_S &= y_i, && \forall i \in [n] \\ z_S &\geq 0, && \forall S \subseteq [n] \end{aligned}$$

Then, H is a convex function.

Proof. Consider two vectors $y^*, y^{**} \in [0,1]^n$. Pick an arbitrary $\lambda \in [0,1]$. We need to show that

$$H(\lambda y^* + (1-\lambda)y^{**}) \le \lambda H(y^*) + (1-\lambda)H(y^{**}).$$

Consider the optimal LP solutions z^* and z^{**} for y^* and y^{**}. Then, by the definition of H,

$$H(y^*) = \sum_{S\subseteq[n]} \big(\sum_{i\in S} d_i\big)^q z^*_S \;\text{ and }\; H(y^{**}) = \sum_{S\subseteq[n]} \big(\sum_{i\in S} d_i\big)^q z^{**}_S.$$

Observe, that $\lambda z^* + (1-\lambda)z^{**}$ is a feasible solution for $\lambda y^* + (1-\lambda)y^{**}$ (since all LP constraints are linear). Hence, $H(\lambda y^* + (1-\lambda)y^{**})$ is at most the LP cost of $\lambda z^* + (1-\lambda)z^{**}$, which equals

$$\sum_{S\subseteq[n]} \Big(\sum_{i\in S} d_i\Big)^q (\lambda z^*_S + (1-\lambda) z^{**}_S) = \lambda H(y^*) + (1-\lambda)H(y^{**}).$$

□

D Discretization

In this section, we show how to discretize values d_{ij}. We assume that functions f_j satisfy the following conditions:

1. All f_j are convex increasing nonnegative functions computable in polynomial time.
2. For every j, $f_j(0) = 0$.
3. For every j and $t \in \big[0, \sum_i d_{ij}\big]$, $t(\log f_j(t))' \le P$.

We first find an approximate value of the optimal solution $\widetilde{OPT} \ge OPT$ and then apply Theorem D.1 (see below). Note that if the gap $\widetilde{OPT}/OPT$ is polynomially bounded, then we can pick ε such that the optimal value OPT' of the discretized problem is at most $(1+\varepsilon')OPT$. We pick such $\widetilde{OPT}$ either by using the binary search or by enumerating all powers of 2 in the range $[\min_{ij} f_j(d_{ij}), \sum_j f_j(\sum_i d_{ij})]$.

Theorem D.1. *There exists a polynomial-time algorithm that given an instance of the integer program (3)-(5) satisfying conditions (1) and (2) above, an upper bound $\widetilde{OPT}$ on the cost of the optimal solution OPT, and $\varepsilon > 0$, returns a new set of coefficients d'_{ij}, numbers $\delta_j > 0$, and an extra set of constraints $y_i = 0$ for $i \in \mathcal{I}$ such that each d'_{ij} is a multiple of δ_j; d'_{ij}/δ_j is an integer polynomially bounded in n, $1/\varepsilon$ and P (see item 3 above) such that the following two properties are satisfied.*

1. *The cost of the optimal solution for the new problem is at most the cost of the original problem:*

$$OPT' \le OPT.$$

2. *For every feasible solution of the new problem $y \in \mathcal{P} \cap \{0,1\}^n$ satisfying $y_i = 0$ for $i \in \mathcal{I}$, we have*

$$\sum_{j\in[k]} f_j\Big(\sum_{i\in[n]} d_{ij} y_i\Big) \le (1+\varepsilon) \sum_{j\in[k]} f_j\Big(\sum_{i\in[n]} d'_{ij} y_i\Big) + \varepsilon\widetilde{OPT}.$$

Particularly,

$$OPT \le (1+\varepsilon)OPT' + \varepsilon\widetilde{OPT} \le OPT' + 2\varepsilon\widetilde{OPT}.$$

Proof. The proof is fairly standard: We round all d_{ij} to be multiples of δ_j. Then we show that if δ_j's are sufficiently small, then the introduced rounding error is at most $2\varepsilon\widetilde{OPT}$. The details are below.

We algorithm finds the set $\mathcal{I} = \{i : f_j(d_{ij}) > \widetilde{OPT} \text{ for some } j\}$. If $i \in \mathcal{I}$, then y_i must be equal to 0 in every optimal solution, because otherwise, $OPT \geq f_j(d_{ij}) > \widetilde{OPT}$. Thus, for all $i \in \mathcal{I}$, we set y_i and all d'_{ij}'s to be 0. Then, we let

$$\eta = \min\big(\frac{\varepsilon}{4P}, \frac{1}{2}\big); \quad t_j = \max_{i \notin \mathcal{I}} d_{ij}; \quad \delta_j = \frac{\varepsilon \eta t_j}{kn}.$$

We round down all d_{ij}'s to be multiples of δ_j. Denote the rounded values by d'_{ij}. This is our new instance.

It is clear that d'_{ij} are multiples of δ_j. Since $t_j \geq d'_{ij}$ for all i and j, we have $d'_{ij}/\delta_j \leq \frac{kn}{\varepsilon\eta}$ and therefore, all d'_{ij}/δ_j are polynomially bounded. Since $d'_{ij} \leq d_{ij}$ and f_j are monotone functions, we have for every y,

$$\sum_{j\in[k]} f_j\Big(\sum_{i\in[n]} d'_{ij} y_i\Big) \leq \sum_{j\in[k]} f_j\Big(\sum_{i\in[n]} d_{ij} y_i\Big).$$

As we observed earlier if y^* is the optimal solution to the original problem, then $y^*_i = 0$ for $i \in \mathcal{I}$, hence y^* is a feasible solution to the new problem. Consequently,

$$OPT' \leq OPT.$$

We now need to verify that

$$\sum_{j\in[k]} f_j\Big(\sum_{i\in[n]} d_{ij} y_i\Big) - \sum_{j\in[k]} f_j\Big(\sum_{i\in[n]} d'_{ij} y_i\Big) \leq \varepsilon\widetilde{OPT} + \varepsilon \sum_{j\in[k]} f_j\Big(\sum_{i\in[n]} d'_{ij} y_i\Big).$$

We prove that for every j,

$$f_j\Big(\sum_{i\in[n]} d_{ij} y_i\Big) - f_j\Big(\sum_{i\in[n]} d'_{ij} y_i\Big) \leq \frac{\varepsilon}{k}\widetilde{OPT} + \varepsilon f_j\Big(\sum_{i\in[n]} d'_{ij} y_i\Big). \tag{43}$$

Consider two cases.

I. If $\sum_{i\in[n]} d_{ij} y_i < \varepsilon t_j/k$, then

$$f_j\Big(\sum_{i\in[n]} d_{ij} y_i\Big) \leq \varepsilon f_j(t_j)/k \leq \varepsilon\widetilde{OPT}/k,$$

since f is a convex function, $f(0) = 0$, and $f_j(t_j) = \max_{i\notin\mathcal{I}} f_j(d_{ij}) \leq \widetilde{OPT}$. Hence, inequality (43) holds.

II. Now assume that $\sum_{i\in[n]} d_{ij} y_i \geq \varepsilon t_j/k$. Observe, that

$$\sum_{i\in[n]} d_{ij} y_i - \sum_{i\in[n]} d'_{ij} y_i \leq n\delta_j = \varepsilon\eta t_j/k \leq \eta \sum_{i\in[n]} d_{ij} y_i.$$

Hence (using that $\eta \in [0, 1/2]$ and thus $1/(1-\eta) \leq 1 + 2\eta$),

$$\sum_{i\in[n]} d_{ij} y_i \leq \frac{1}{1-\eta} \sum_{i\in[n]} d'_{ij} y_i \leq (1+2\eta) \sum_{i\in[n]} d'_{ij} y_i.$$

We now use Claim D.2 (see below) with

$$t = \sum_{i\in[n]} d'_{ij} y_i; \quad \text{and} \quad \eta' = \frac{\sum_{i\in[n]} d_{ij} y_i}{\sum_{i\in[n]} d'_{ij} y_i} - 1 \le 2\eta \le \frac{\varepsilon}{2P}.$$

We get

$$f_j\Big(\sum_{i\in[n]} d_{ij} y_i\Big) - f_j\Big(\sum_{i\in[n]} d'_{ij} y_i\Big) \le \varepsilon f_j\Big(\sum_{i\in[n]} d'_{ij} y_i\Big).$$

This finishes the proof. It only remains to prove Claim D.2. □

Claim D.2. *Suppose that f is a nonnegative monotonically increasing function such that for every $t \in (0, T]$, $t(\log f(t))' \le P$. Let $\varepsilon \in [0, 1]$ and $\eta' \le \varepsilon/(2P)$. Then, for $t \in (0, T/(1+\eta'))$,*

$$f((1+\eta')t) - f(t) \le \varepsilon f(t).$$

Proof. Write,

$$\begin{aligned}\log f((1+\eta')t) - \log f(t) &= \int_t^{(1+\eta')t} (\log f(s))' ds \le \int_t^{(1+\eta')t} \frac{P}{s} ds \\ &= P \log \frac{(1+\eta')t}{t} = P \log(1+\eta') \le P\eta'.\end{aligned}$$

Thus,

$$\frac{f((1+\eta')t)}{f(t)} \le e^{\eta' P} \le e^{\varepsilon/2},$$

and (since $e^{\varepsilon/2} - 1 \le \varepsilon$ for $\varepsilon \in [0, 1]$)

$$f((1+\eta')t) - f(t) \le (e^{\varepsilon/2} - 1) f(t) \le \varepsilon f(t).$$

□

E Lemma E.1

Lemma E.1. *Suppose that a sequence of integer random variables S^n converges in distribution to an integer random variable Z. Then, for every nonnegative function φ with a finite expectation $\mathbb{E}[\varphi(Z)]$, we have*

$$\sup_n \mathbb{E}[\varphi(S^n)] \ge \mathbb{E}[\varphi(Z)].$$

Proof. Let $\varphi_t(x) = \varphi(x)$ for $x \le t$, and $\varphi_t(x) = 0$, otherwise. The function φ is nonzero on finitely many integral points. Hence, $\lim_{n\to\infty} \mathbb{E}[\varphi_t(S^n)] = \mathbb{E}[\varphi_t(Z)]$ for every fixed t. Observe that $\mathbb{E}[\varphi(S^n)] \ge \mathbb{E}[\varphi_t(S^n)]$, since $\varphi(x) \ge \varphi_t(x)$ for all x. On the other hand, $\lim_{t\to\infty} \mathbb{E}[\varphi_t(Z)] = \mathbb{E}[\varphi(Z)]$. Hence,

$$\sup_n \mathbb{E}[\varphi(S^n)] \ge \sup_n \sup_t \mathbb{E}[\varphi_t(S^n)] = \sup_t \sup_n \mathbb{E}[\varphi_t(S^n)] \ge \sup_t \lim_{n\to\infty} \mathbb{E}[\varphi_t(S^n)] = \sup_t \mathbb{E}[\varphi_t(Z)] = \mathbb{E}[\varphi(Z)].$$

□

F Convex Order

For completeness, we give proofs of Lemma 5.5 and Lemma 5.6 in this section. The reader may also find slightly different proofs of these lemmas in the book of Shaked and Shanthikumar [32].

Lemma F.1. *Consider three independent random variables X, Y, and Z. If $X \leq_{cx} Y$, then $Z + X \leq_{cx} Z + Y$.*

Proof. For every convex function $\varphi : \mathbb{R} \to \mathbb{R}$, we have

$$\mathbb{E}\left[\varphi(X+Z)\right] = \mathbb{E}\mathbb{E}[\varphi(X+Z) \mid Z] \leq \mathbb{E}\mathbb{E}[\varphi(Y+Z) \mid Z] = \mathbb{E}\left[\varphi(Y+Z)\right].$$

The inequality above holds, since for any fixed $Z = z$, the function $x \mapsto \varphi(x+z)$ is convex. □

Proof of Lemma 5.5. Using Lemma F.1, we replace X_i's with Y_i's in the sum $X_1 + \cdots + X_n$ one by one: For every k, let

$$S_k = \sum_{i=1}^{k} X_i + \sum_{j=k+1}^{n} Y_i.$$

Then, by Lemma F.1,

$$S_k = \Big(\sum_{i=1}^{k-1} X_i + \sum_{j=k+1}^{n} Y_j\Big) + X_k \leq_{cx} \Big(\sum_{i=1}^{k-1} X_i + \sum_{j=k+1}^{n} Y_j\Big) + Y_k = S_{k-1}.$$

We have

$$\sum_{i=1}^{n} X_i = S_n \leq_{cx} S_{n-1} \leq_{cx} \ldots \leq_{cx} S_1 = \sum_{i=1}^{n} Y_i.$$

This finishes the proof. □

Proof of Lemma 5.6. Consider an arbitrary convex function φ. Let $A = \sum_{i=1}^{n} a_i$ and $\tilde{\varphi}(x) = \varphi(Ax)$. Since $\tilde{\varphi}$ is a convex function, we have

$$\tilde{\varphi}\Big(\frac{1}{A}\sum_{i=1}^{n} a_i X_i\Big) \leq \frac{1}{A}\sum_{i=1}^{n} a_i \tilde{\varphi}(X_i),$$

and $\mathbb{E}\left[\tilde{\varphi}(X_i)\right] \leq \mathbb{E}\left[\tilde{\varphi}(Y)\right]$. Hence,

$$\mathbb{E}\left[\tilde{\varphi}\Big(\frac{1}{A}\sum_{i=1}^{n} a_i X_i\Big)\right] \leq \mathbb{E}\left[\frac{1}{A}\sum_{i=1}^{n} a_i \tilde{\varphi}(X_i)\right] = \frac{1}{A}\sum_{i=1}^{n} a_i \mathbb{E}\left[\tilde{\varphi}(X_i)\right] \leq_{cx} \frac{1}{A}\sum_{i=1}^{n} a_i \mathbb{E}\left[\tilde{\varphi}(Y)\right] = \mathbb{E}\left[\tilde{\varphi}(Y)\right].$$

Substituting $\tilde{\varphi}(x) = \varphi(Ax)$, we get

$$\mathbb{E}\left[\varphi\Big(\sum_{i=1}^{n} a_i X_i\Big)\right] \leq \mathbb{E}\left[\tilde{\varphi}\Big(\sum_{i=1}^{n} a_i Y\Big)\right].$$

□

G Continuous Random Variables

In Section 5, we proved Theorem 5.3 for discrete random variables. In this section, we extend this result to arbitrary random variables. Consider two sequences of nonnegative random variables $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_n$ satisfying conditions of Theorem 5.3. Fix an arbitrary convex function $\varphi : \mathbb{R} \to \mathbb{R}$. We need to show that

$$\mathbb{E}[\varphi(X_1 + \cdots + X_n)] \leq \mathbb{E}[\varphi(P \cdot (Y_1 + \cdots + Y_n))] \tag{44}$$

assuming that both expectations exist. Observe that φ can be represented as the sum of two convex functions: a monotonically non-decreasing convex function $\varphi^{\uparrow}$ and monotonically non-increasing convex function $\varphi^{\downarrow}$. It suffices to show that

$$\mathbb{E}[\varphi^{\uparrow}(X_1 + \cdots + X_n)] \leq \mathbb{E}[\varphi^{\uparrow}(P \cdot (Y_1 + \cdots + Y_n))]; \tag{45}$$

$$\mathbb{E}[\varphi^{\downarrow}(X_1 + \cdots + X_n)] \leq \mathbb{E}[\varphi^{\downarrow}(P \cdot (Y_1 + \cdots + Y_n))]. \tag{46}$$

Note that if the expectations in (44) exist than the expectations in (45) and (46) also exist. We now prove inequality (45). The proof of (46) is almost the same. For natural M, define a function g_M as follows:

$$g_M(x) = \begin{cases} \frac{\lfloor Mx \rfloor}{M}, & \text{if } x \leq M; \\ 0, & \text{otherwise.} \end{cases}$$

The function g_M truncates x at the level M and then rounds x down to the nearest multiple of $1/M$. Observe, that $\lim_{M\to\infty} g_M(x) = x$ for every x. Hence, $(g_M(X_1) + \cdots + g_M(X_n))$ converges a.s. to $X_1 + \cdots + X_n$ as $M \to \infty$; and $P(g_M(Y_1) + \cdots + g_M(Y_n))$ converges a.s. to $P(Y_1 + \cdots + Y_n)$ as $M \to \infty$. Since, $\varphi^{\uparrow}$ is a continuous function $\varphi^{\uparrow}(g_M(X_1) + \cdots + g_M(X_n)) \xrightarrow{a.s.} \varphi^{\uparrow}(X_1 + \cdots + X_n)$ and $\varphi^{\uparrow}(P(g_M(Y_1) + \cdots + g_M(Y_n))) \xrightarrow{a.s.} \varphi^{\uparrow}(P(Y_1 + \cdots + Y_n))$. Notice that $g_M(x) \leq x$ for all x. Hence, $\varphi^{\uparrow}(g_M(X_1) + \cdots + g_M(X_n)) \leq \varphi^{\uparrow}(X_1 + \cdots + X_n)$ and $\varphi^{\uparrow}(P(g_M(Y_1) + \cdots + g_M(Y_n))) \leq \varphi^{\uparrow}(P(Y_1 + \cdots + Y_n))$. By Lebesgue's dominated convergence theorem, we get

$$\lim_{M\to\infty} \mathbb{E}[\varphi^{\uparrow}(g_M(X_1) + \cdots + g_M(X_n))] = \mathbb{E}[\varphi^{\uparrow}(X_1 + \cdots + X_n)]; \tag{47}$$

$$\lim_{M\to\infty} \mathbb{E}[\varphi^{\uparrow}(P(g_M(Y_1) + \cdots + g_M(Y_n)))] = \mathbb{E}[\varphi^{\uparrow}(P(Y_1 + \cdots + Y_n))]. \tag{48}$$

For every fixed M, the left hand side of (47) is upper bounded by the left hand side of (48), because random variables $g_M(X_1), \ldots, g_M(X_n)$ and $g_M(Y_1), \ldots, g_M(Y_n)$ are discrete and satisfy the conditions of Theorem 5.3. Hence, the right hand side of (47) is upper bounded by the right hand side of (48). This proves inequality (45) and concludes the proof of Theorem 5.3 for arbitrary random variables.

H Figures

$q =$	1	1.25	1.5	1.75	2
$A_q =$	1	1.163	1.373	1.645	2

Figure 1: The values of A_q for some $q \in [1, 2]$. All values are rounded up to three decimal places.

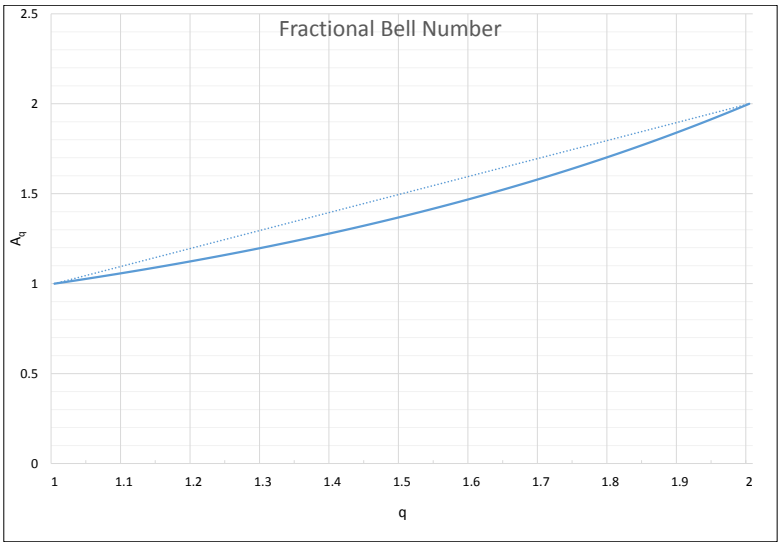

Figure 2: Graph of A_q for $q \in [1, 2]$. Note that the function $q \mapsto A_q$ is convex; $A_1 = 1$ and $A_2 = 2$. Thus, $A_q \leq q$ for $q \in [1, 2]$.

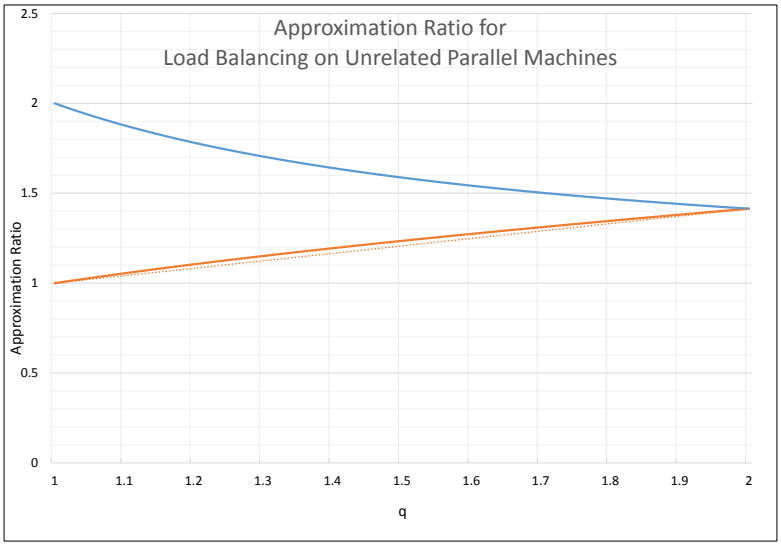

Figure 3: Approximation factor of our algorithm for Load Balancing on Unrelated Parallel Machines – $A_q^{1/q}$ – is plotted in red (below). Approximation factor of the algorithm due to Kumar, Marathe, Parthasarathy and Srinivasan [22] – $2^{1/q}$ – is plotted in blue (above). The function $A_q^{1/q}$ can be well approximated by the linear function $1 + (\sqrt{2} - 1)(q - 1)$ in the interval $q \in [1, 2]$.